



CE Workgroup

Status of Embedded Linux September 2018

Tim Bird

Architecture Group Chair


LF Core Embedded Linux Project



CE Workgroup

Nature of this talk...

- Quick overview of lots of embedded topics
- A springboard for further research
 - If you see something interesting, you have a link or something to search for
- Not comprehensive!
 - Just stuff that I saw





CE Workgroup

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Kernel Versions



- Linux v4.14 – 12 Nov 2017 – 70 days
- Linux v4.15 – 28 Jan 2018 – 77 days
 - Included Spectre and Meltdown fixes
- Linux v4.16 – 1 Apr 2018 – 63 days
- Linux v4.17 – 3 Jun 2018 – 63 days
- Linux v4.18 – 12 Aug 2018 – 70 days
- Now on Linux v4.19-rc5
 - **Author: Greg Kroah-Hartman**
 - Expect 4.19 on Oct 21



CE Workgroup

Linux 4.14

- New kernel stack unwinder (ORC) for x86_64
 - Better unwinding via kernel-specific out-of-band structure (for every kernel PC address)
 - See <https://lwn.net/Articles/728339/>
- zstd compression for btrfs and squashfs
- Better cpufreq coordination with SMP





CE Workgroup

Linux 4.15

- Cramfs supports mapping persistent memory
 - Can use for XIP
- AMD display core system accepted
- Device tree compiler has support for overlays
- RISC-V support
- Spectre/Meltdown mitigations
 - KPTI
 - retpolines



CE Workgroup

Linux 4.16

- Initial support for the Jailhouse hypervisor
- eBPF support for functions
- arm64 mitigations for Spectre and Meltdown
- More Spectre mitigations (general)
 - `array_index_nospec()`
- High resolution timers now have two modes, to allow them to be run in software interrupt context



CE Workgroup

Linux 4.16 – cont.

- F2FS miscellaneous improvements
- Slimbus and Soundwire sub-systems added
 - These are MIPI audio bus standards
- Flex and Bison are required for kernel build



CE Workgroup

Linux 4.17

- 8 old architectures dropped
 - Blackfin, CRIS, FRV, M32R, Metag, MN10300, Score, Tile
 - Removes about 460K lines of code
 - Only 3rd time ever that a kernel release has shrunk
- Rework of kernel idle loop
- Finished full in-kernel TLS protocol support
- Improved CPU load estimation



CE Workgroup

Improved CPU load estimation

- Is a modification of the per-entity load-tracking (PELT) mechanism
 - PELT decays the load information about processes too quickly
 - New estimator avoids this
- Load estimation can clamp more quickly
- Good for mobile and embedded
- Adds 1% scheduling overhead
 - Requires setting SCHED_UTILEST scheduler feature bit
- See <https://lwn.net/Articles/741171/>



CE Workgroup

Linux 4.17 – cont.

- A formal kernel memory-ordering model
 - With tests for formal proofs of adherence
 - See <https://lwn.net/Articles/718628/>
- Kernel build now requires gcc 4.5 or later on x86
 - This is a problem for some architectures where gcc has dropped support
 - But there's a workaround:
 - <https://lists.linuxfoundation.org/pipermail/ksummit-discuss/2018-September/005554.html>
- Changes to x86 system call implementation



CE Workgroup

Linux 4.18

- power domains now support active state management
 - Instead of enable/disable, can now handle different idle states (continuum of operation)
 - <https://lwn.net/Articles/744047/>
- fscrypt supports Speck128 and Speck256 ciphers
 - Somewhat controversial ciphers
 - Enables encryption for lowest-end devices
- bpfILTER user mode helper system
 - <https://lwn.net/Articles/755919/>



CE Workgroup

bpfilter user mode helper




- Complicated mechanism to:
 - compile user space code
 - from the Linux source tree
 - bundled in a kernel loadable module file (.ko)
 - with execution initiated from kernel space
- Used to provide support for backwards compatibility with netfilter configuration protocol
 - user-space 'program' to compile netfilter config protocol into bpf pseudo-code
- As a mechanism, may lead to all kinds of crazy stuff



CE Workgroup

Linux 4.18 (cont.)

- Support for Qualcomm Snapdragon 845
 - Use in high-end mobile devices
 - Support is incomplete, but it's a start





CE Workgroup

Linux 4.19 (expected)

- L1TF mitigation (a variant of meltdown)
 - <https://lwn.net/Articles/762570/>
 - The fun continues...
- time-based packet transmission
 - Allows a program to schedule data for transmission in the future
 - <https://lwn.net/Articles/748879/>
- EROFS – enhanced read-only filesystem
 - High-performance
 - Good for certain embedded situations
 - (in staging)




CE Workgroup

Contributor stats for 4.18

- 12,879 changes sets (as of rc6)
- 553K lines added, 652K lines removed
- 1668 developers
 - 226 made their first contribution

First-time contributors per release cycle



Source: <https://lwn.net/Articles/760690>



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff

Resources



CE Workgroup

Bootup Time

- Nothing new, here is older stuff...
- Analyze_boot tool – new in in 4.12
- Some good previous talks:
 - ELCE 2017 - *A Pragmatic Guide to Boot-Time Optimization* by Chris Simmonds
 - ELCE 2014 - *12 Lessons Learnt in Boot Time Reduction* by Andrew Murray



CE Workgroup

Device Tree

- **Nothing new, here is older stuff...**
- **Device Tree validation**
 - Schema for binding language, validator for bindings and for device tree data
 - New proposal for device tree validation by Pantellis and Grant Likely
 - Implementation is in-progress
- **Updated Device Tree specification**
 - Want to update material and make it more available
- **Overlays**
 - Device tree compiler has support for overlays



CE Workgroup

File Systems




- **F2FS**
 - Miscellaneous fixups (4.17)
 - Lost & Found support
 - Better tuning for low_end devices
 - See https://www.phoronix.com/scan.php?page=news_item&px=F2FS-Lost-Found
 - Support for disk quotes (4.13, 4.15)
- **BTRFS and Squashfs support for zstd compression (4.14)**
 - Faster and smaller compression/decompression
 - <https://clearlinux.org/blogs/linux-os-data-compression-options-comparing-behavior>
 - See https://www.phoronix.com/scan.php?page=news_item&px=Linux-4.14-Zstd-Pull



CE Workgroup

Graphics

- Working on support for virtual reality
 - LCA 2018 *Driving Virtual Reality from Linux* - Keith Packard





CE Workgroup

GPU drivers

- ELC 2018 *Progress in the Embedded GPU Ecosystem* – by Robert Foss
 - *Watch the video – the slides don't have enough text*
 - Nvidia, Intel, AMD, Broadcom, Qualcomm, Vivante have upstream support
 - Of varying quality
 - ARM – some stuff happening recently with Mali T series, but not upstreamed yet.



Networking

- Time Sensitive Networking
 - ELC 2018 *The Road Towards a Linux TSN Infrastructure* – Jesus Sanchez-Palencia
 - ELCE 2017 *Deterministic Networking for Real-Time Systems (Using TSN)* – by Henrik Austad
 - so_txtime option for high-resolution transmit time
 - IEEE deterministic networking (DetNet) working group
 - Lots of standards
- Time-based packet transmission
 - Allows a program to schedule data for transmission in the future
 - <https://lwn.net/Articles/748879/>
- Bluetooth 5 – supported



CE Workgroup

Power management

- Rework kernel idle loop (in 4.17)
 - Prevent CPUs from spending too much time in shallow idle states
 - Reduces idle power on some systems by 10% or more
 - See
 - <https://www.phoronix.com/scan.php?page=article&item=linux-417-power>
 - <https://lkml.org/lkml/2018/4/11/337>
- power domain state management (4.18)
 - Instead of enable/disable, can now handle different idle states (continuum of operation)
 - <https://lwn.net/Articles/744047/>



Power Management

- Presentations:
 - ELC 2018 *An Unbiased Look at the Energy Aware Scheduler (EAS)* – by Vital Wool
 - Qualcomm has their own big.LITTLE scheduler (QHMP)
 - QHMP does better than EAS in some regards
 - But cannot be mainlined (code is messy)
 - Want to use features of QHMP in EAS, which still has shortcomings



CE Workgroup

Real Time

- RT-Preempt patches give good real-time performance
- RT-Preempt patch still out of tree
 - What's left:
 - Hotplug locking
 - Timer wheel rework
 - dentry cache locking
 - Lots work goes into maintaining RT trees out-of-mainline
 - Don't support every kernel release
 - Focused on supporting kernel LTS releases



Real Time (cont.)

- Presentations
 - ELC 2018 *Steering Xenomai into the Real-Time Linux Future* – Jan Kiska
 - ELC 2018 *Not Really, but Kind of Real Time Linux* – Sandra Capri
 - Discusses how much RT performance you can get, without Preempt-RT patches
 - ELC 2018 *Preempt-RT Raspberry Pi Linux* – Tiejun Chen
 - Demonstrates that Preempt-RT is very effective on Raspberry Pi
 - ELC 2018 *Maintaining a Real Time Stable Kernel* – by Steven Rostedt



CE Workgroup

Security (review)

- Spectre and Meltdown
 - Break security via side-channel timing attacks using speculative execution
 - Variants 1, 2 (Spectre), and 3 (Meltdown)
- Is a family of vulnerabilities related to speculative execution
 - Many modern processors vulnerable
 - Many embedded processors not affected
- Very severe problem:
 - Can read data you're not supposed to
 - Vulnerability has existed for 20 years!
 - Cannot be fixed with CPU firmware updates
 - Mitigations are expensive



CE Workgroup

Security

- New Spectre variants
 - Variant 3a – Rogue System Register Read
 - Variant 4 – Speculative Store Bypass
- New Meltdown variant
 - L1TF – L1 Cache terminal fault vulnerability
- No surprise
 - We were expecting new variations of speculative execution vulnerabilities to be discovered
- Fixes are:
 - More microcode updates for Intel processors
 - Kernel patches to use new speculative execution control flags
- See <https://lwn.net/Articles/755114/>



CE Workgroup

Security Presentations

- ELC 2018 *Secure Boot from A to Z* – by Quentin Schulz and Mylune Josserand
 - Overview of secure boot techniques and issues
- ELCE 2017 *Security Features for UBIFS* – by Richard Weinberger



CE Workgroup

System Size




- No new kernel features
- Presentations
 - ELC 2018 *Poky-tiny and Beyond, or Trying to the put Yocto in Yocto Project* – by Scott Murray
 - Gives status of poky-tiny project, available for Yocto Project
 - ELC 2018 BoF: *Embedded Linux Size* – By Michael Opdenacker
 - Great overview of reduction techniques and status
 - Toybox and musl (smaller libc) are worth looking at
 - Long list of things that can be worked on



CE Workgroup

Testing

- Kselftest
- Fuego
- Kernelci.org
- LKFT
- Work to make 'next' more testable





CE Workgroup

Kselftest

- Nothing new, here is older stuff...
- Unit test system inside kernel source tree
- Recent work:
 - -silent option, to reduce output clutter
 - Support for O= option, to build outside source directory
 - Lots more regression tests (preferred place for syscall compatibility/regression tests (over LTP))
 - Converting to TAP (Test Anything Protocol) for test output (started in 4.13)
- See <https://lwn.net/Articles/737893/>



CE Workgroup

Fuego

- Test Framework for collaborating on tests and test infrastructure for Linux
- v1.3 released May 2018
 - More report output formats (rst, csv, excel, html, pdf)
 - Hardware board control
 - Test phase execution
- Tests being added on a consistent basis
 - 18 new tests in 1.3 release (some are self-tests)
 - 7 are realtime tests
- Presentation:
 - Japan Jamboree 63: *Fuego Status and Roadmap December 2017* – by Tim Bird



CE Workgroup

Kernelci.org

- Does continuous build/boot testing of kernel
 - Builds 126 trees continuously, then reports any errors
- Working on creating a project in Linux Foundation (more later)



CE Workgroup

LKFT

- Linux Kernel Functional Testing
 - Relatively new Linaro kernel testing effort
 - Focused on Functional testing (as opposed to build/boot testing)
 - Focused on embedded devices
- Presentation:
 - ELC 2018 *Keeping Up With LTS: Linux Kernel Functional Testing (LKFT) on Devices* – Thomas Gall



CE Workgroup

Making 'next' more testable

- Linux-next is the integration tree used during the kernel release cycle
- It's hard to test, because things break a lot
 - Automated testing doesn't work
- Stephen Rothwell (the 'next' maintainer) created a 'fixes' branch
- Isolates fixes intended for next release, from other code being integrated into 'next'
 - Should not break automated testing rigs as much
- Result: fixes will get more testing in 'next'



CE Workgroup

Toolchains

- gcc 8
 - Major effort on usability improvements
 - Provides much better messages for some errors
 - Shows fix-it hints
 - Shows what to change to fix the error
 - Can be automatically processed
 - Detects missing include files, saying which files are needed
 - See <https://lwn.net/Articles/749450/>, and
 - <https://developers.redhat.com/blog/2018/03/15/gcc-8-usability-improvements/>



Tracing





- Nothing new, here is older stuff...
- Dynamic function tracing events
 - Ability to create a tracepoint for a function at runtime
 - Goal is to avoid having a tracepoint become part of kernel ABI
 - Is work-in-progress
 - See <https://lwn.net/Articles/747256>
- Presentations:
 - ELC 2017 *Dynamic Tracing Tools on ARM/AArch64 Platform: Updates and Challenges* - by Hiroyuki Ishii
 - Great overview of Linux tracing capabilities and programs



CE Workgroup

Miscellaneous

- Year 2038 work
- Git protocol version 2
- Android kernel status





CE Workgroup

Year 2038 work

- Status update:
 - Lots of small driver fixes in 4.16
 - Changes to system call entry points for timekeeping related syscalls
 - Patches for structures with new 64-bit timestamps have been submitted
 - Still need more work converting the VFS layer
 - Lots of stuff intended to land in 4.18
- See <https://www.mail-archive.com/linux-kernel@vger.kernel.org/msg1674216.html>



CE Workgroup

New git protocol (version 2)

- 3x performance improvement for no-op fetches on repositories containing 500k references.
- 8x reduction of overhead bytes sent from server
 - Due to filtering references to those the client expressed interest in
- Worked on by Google
 - See <https://opensource.googleblog.com/2018/05/introducing-git-protocol-version-2.html>



CE Workgroup

Android kernel status

- Progress being made
- diff from 4.14 Android and LTS
 - 432 files, 41K changes
 - sdcard, netfilter, Energy Aware Scheduling, USB gadgets
- Linaro doing android mainline tracking
 - Test Android-common patches on latest mainline Linux
- ELC 2018 *Android Common Kernel and Out of Mainline Patchset Status* – by Amil Pundar



CE Workgroup

Outline

Kernel Versions

Technology Areas

CE Workgroup Projects

Other Stuff



Resources



CE Workgroup

Projects and initiatives

- Shared Embedded Distribution
- LTSI
- Automated Testing
 - Fuego
 - Automated Test Standards
- eLinux wiki





CE Workgroup

Shared Embedded Distribution

- **Goals**
 - Create an industry-supported distribution of embedded Linux
 - Main goal is very long term support (15 years)
- **Status**
 - Working on building Debian with Yocto Project
 - 3 projects - meta-debian, isar and elbe wish to collaborate and combine their yocto recipes into a single layer.
- **Next steps**
 - Continued integration of Debian-based build and packaging systems



CE Workgroup

Long Term Support Initiative

- LTSI 4.9 is current LTSI kernel
 - Work is in progress on next release 4.14
 - 4.14.70-ltsi-rc1 is now available
- Most of industry is using LTS or LTSI
- Using upstream-first policy for patches
- Security fixes are very important
- Presentation:
 - *ELCE 2017 Using Long Term Stable Kernel for the Embedded Products* – by Tsugikazu Shibata



CE Workgroup

Long Term Stable Releases

Version	Maintainer	Released	Projected EOL
3.16	Ben Hutchings	2014-08-03	Apr, 2020
4.4	Greg Kroah-Hartman	2016-01-10	Feb, 2022
4.9	Greg Kroah-Hartman	2016-12-11	Jan, 2023
4.14	Greg Kroah-Hartman	2017-11-12	Jan, 2020
4.19	Greg Kroah-Hartman	2018-??	TBD

Source: <https://www.kernel.org/category/releases.html>



CE Workgroup

Fuego - Linux Test Framework

- CELP funding for Fuego self-test project
- Fuego now has an integrated release test
 - A Fuego job to build the Fuego docker container from scratch, and test it
 - Includes tests of user interface using Selenium and Chromium
 - This adds packages to base Fuego distribution for doing this type of web-based and image-based testing
- Work completed by ProFusion Embedded Systems



CE Workgroup

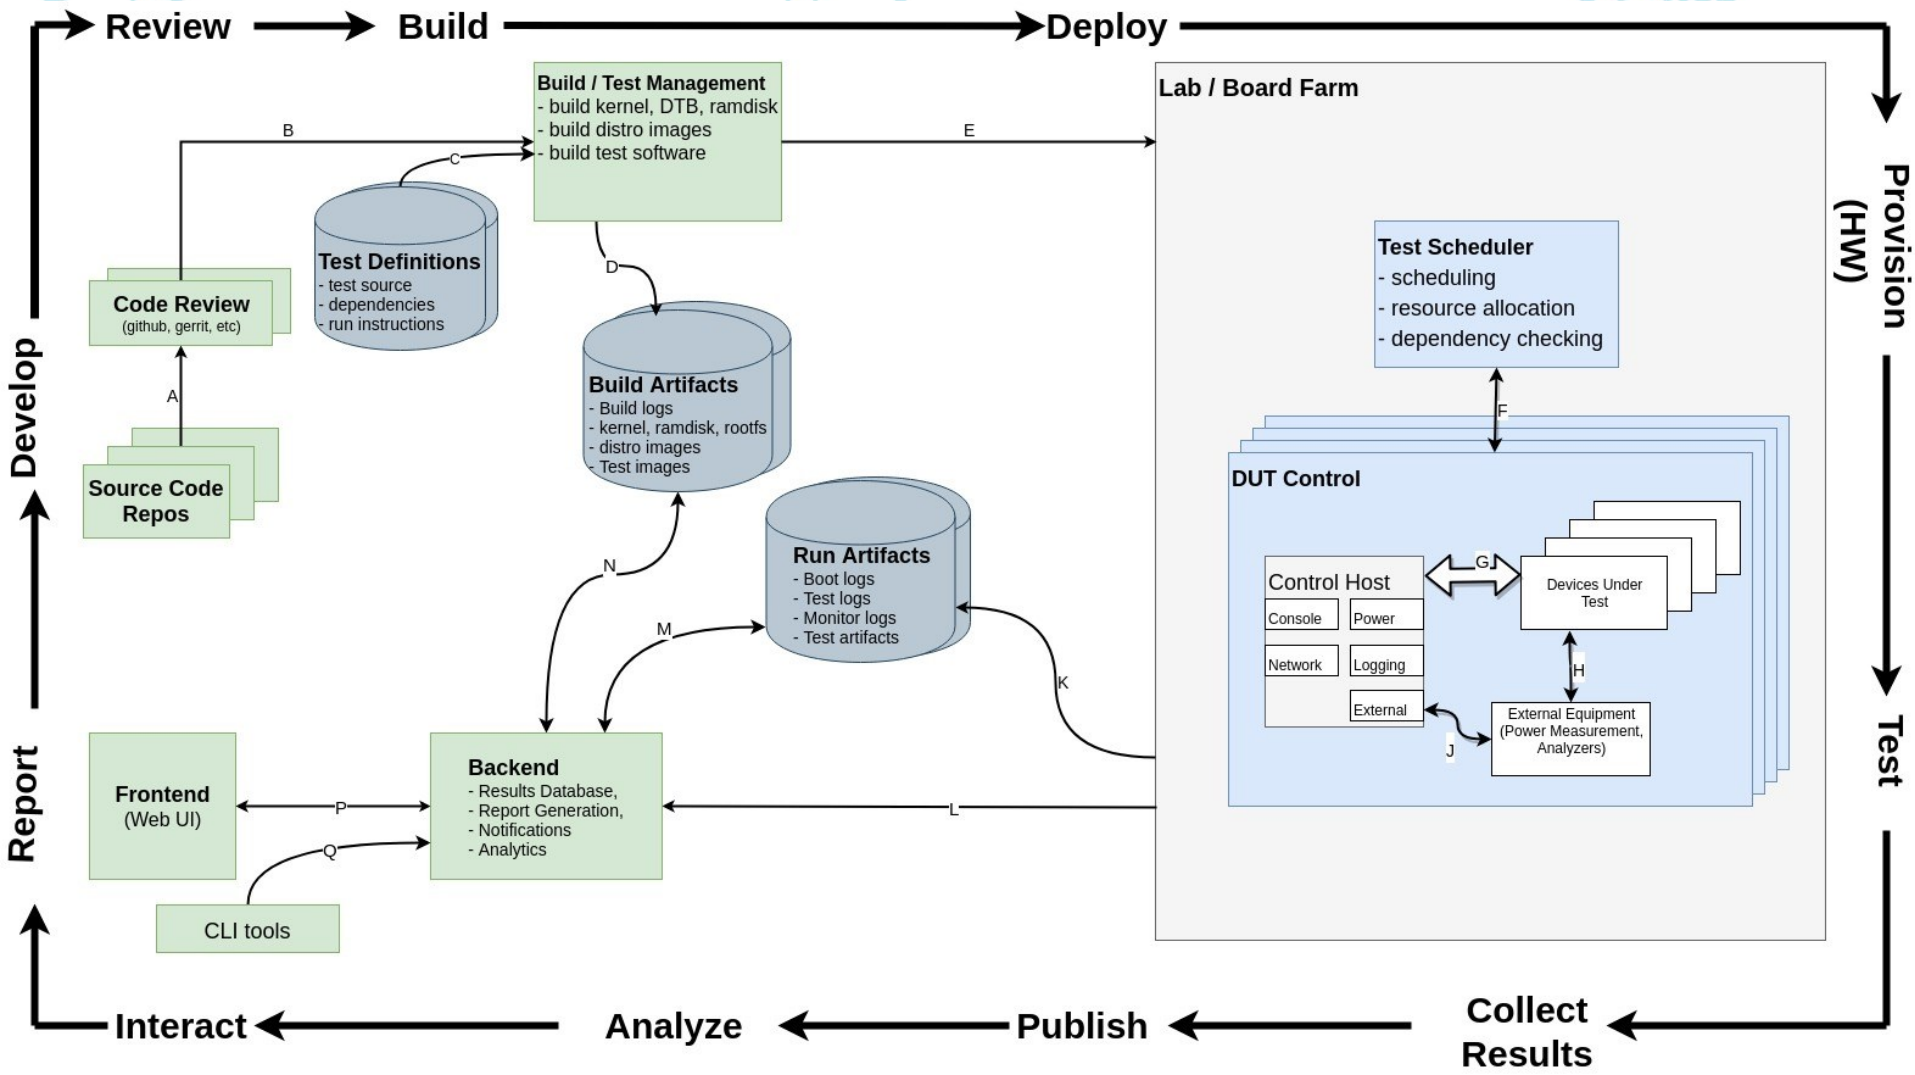
Automated Testing Standards

- CELP is a sponsor of the Automated Testing Summit
 - Summit of test framework architects
 - Goal is to increase collaboration between projects
- Survey of test frameworks
 - See https://elinux.org/Test_Stack_Survey
- Discussions on:
 - automated-testing@yoctoproject.org
- Development of Test Stack model
 - Kevin Hilman (kernelci) and Tim Bird (Fuego)



CE Workgroup

Continuous Integration Loop (high level diagram)





CE Workgroup

eLinux wiki



- <http://elinux.org>
 - Web site dedicated to information for embedded Linux developers
 - The wikipedia of embedded linux!
- Hundreds of pages covering numerous topic areas: bootup time, realtime, security, power management, flash filesystem, toolchain, editors
- **Slides and Videos for 12 years of ELC!!**
- Please use and add to site



CE Workgroup

eLinux wiki

- Recent topics
 - Board farm and automated testing pages
 - Lots of Renesas board information
 - Developer guidelines
 - Community Doc Translation
 - Event pages (ELC, Jamboree, and others)







CE Workgroup

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources









CE Workgroup

Other Stuff

- Community issues
- Trade associations
- Conferences
- Legal issues
- Industry changes








CE Workgroup

Community issues

- Linus takes a break!!
- New Code of Conduct for the kernel community





CE Workgroup

Linus takes a break

- Linus announced on Sept. 16 (4.19-rc4) that he's taking a break
- Made an apology for “flippant attacks in emails that have been both unprofessional and uncalled for”
- Announced that Greg Kroah-Hartman will take his place for 4.19 release cycle
- See
 - <https://lkml.org/lkml/2018/9/16/167>
 - <https://lwn.net/Articles/765108>



CE Workgroup

Linus' break (cont.)

- Linus says he's not burnt out, but will be back
- Needs time to reflect
- Is seeking help on how to behave differently
 - BBC article: Linus Torvalds: "I'll never be cuddly but I can be more polite"
 - <https://www.bbc.com/news/technology-45664640>
- Will be at Maintainer's Summit in Edinburgh (October 22)



CE Workgroup

New Code of Conduct

- Linus also accepted a patch from GregKH that replaces the “code of conflict” with a new “Code of Conduct”
- Based on widely used Contributor Covenant (version 1.4)
- Lots of discussion in community about CoC
 - Concern over new responsibilities for maintainers
 - Uncertainty over enforcement policies
 - <https://lwn.net/Articles/766699/> (likely still paywalled)
 - Ask, and I can forward a link; it will be free-to-view in a week



CE Workgroup

Trade associations

- Linux Foundation
 - Possible creation of KernelCI Testing project
 - KernelCI developers working on getting new hosting
 - Project is underfunded by Linaro
 - Project may expand scope (remains to be seen)



CE Workgroup

Conferences


- Embedded Linux Conference 2018
 - March 12-14, Portland, Oregon, USA
 - See https://elinux.org/ELC_2018_Presentations
 - Did really good at collecting slides and videos
- Japan Jamborees
 - Continuing
- Open Source Summit Japan/Automotive Linux Summit
 - June 20-22, Tokyo, Japan
- ELC Europe 2018
 - October 22-24, Edinburgh, Scotland
- Automated Testing Summit
 - October 25, Edinburgh, Scotland



CE Workgroup

ELCE 2018 topic clusters

- Kernel drivers
 - Camera, Audio, others
- Testing
- Yocto Project
- Security
- Bootloader
- Virtualization
- Realtime
- Networking





CE Workgroup

Legal issues

- McHardy withdraws suit against Geniatech in Germany
- Geniatech fought back, with arguments:
 - Suit scope is too broad (covered all kernel versions, not just ones McHardy had contributed to)
 - Did not show that his commits fulfilled requirements for copyright protection
 - Did not show which of his commits were used by Geniatech
 - McHardy is not following community norms, with regard to GPL revocation terms
 - McHardy is approaching multiple companies for monetary gain



CE Workgroup

Legal issues (cont.)

- McHardy withdrawal – lessons learned:
 - Don't sign the cease-and-desist declaration
 - Ensure GPL compliance
 - Prepare a legal defense strategy
 - Geniatech arguments seem sound, and can be used elsewhere
- Community wants to fight McHardy, but still allow for proper legal enforcement of GPL
- See <https://lwn.net/Articles/752485/>
- Details: <http://laforge.gnumonks.org/blog/20180307-mchardy-gpl/>



CE Workgroup

Industry changes

- Intel selling Wind River (March 2018)
 - <https://www.windriver.com/news/press/pr.html?ID=20982>
 - Not sure what this means for Yocto Project
 - Intel has discontinued Edison, Galileo and Joule
- Microsoft acquires github (June 2018)
 - <https://techcrunch.com/2018/06/04/microsoft-has-acquired-github-for-7-5b-in-microsoft-stock/>



CE Workgroup

Outline

Kernel Versions
Technology Areas
CE Workgroup Projects
Other Stuff
Resources





CE Workgroup

Resources

- LWN.net
 - <http://lwn.net/>
 - If you are not subscribed, please do so
- Kernel Newbies
 - http://kernelnewbies.org/Linux_4.??
- Phoronix
 - <https://www.phoronix.com/>
- eLinux wiki - <http://elinux.org/>
 - Especially <http://elinux.org/Events> for slides and videos



CE Workgroup



Thanks!

